

CLAIMS

What is claimed is:

- 1 1. In a data mining agent executing in a computer system, a method of data
2 mining comprising the steps of:
3 examining a request queue comprising at least one request for data mining
4 processing;
5 determining if the at least one request for data mining processing can be
6 processed;
7 accepting the at least one request for data mining processing if it is
8 determined that the at least one request for data mining processing can be processed;
9 and
10 processing the accepted request for data mining processing in the computer
11 system.
- 1 2. The method of claim 1, wherein the determining step comprises the steps of:
2 determining if an algorithm required to process the at least one request for
3 data mining processing is supported by the computer system;
4 if the algorithm required to process the at least one request for data mining
5 processing is supported, determining whether the computer system is available for
6 additional processing;

7 if the computer system is not available for additional processing, determining
8 whether the computer system will become available for additional processing before
9 other computer systems that might process the at least one request;

10 if the computer system is available for additional processing, or if the
11 computer system will become available for additional processing before other
12 computer systems that might process the at least one request, determining whether
13 the computer system will be able to complete requested processing in an allotted
14 time; and

15 if the computer system will be able to complete the requested processing in
16 the allotted time, determining that the computer system can process the at least one
17 request for data mining processing.

1 3. The method of claim 2, wherein the at least one request for data mining
2 processing comprises data defining at least one algorithm that must be performed in
3 order to perform the requested data mining processing.

1 4. The method of claim 3, wherein there is data defining algorithms that are
2 supported by the computer system.

1 5. The method of claim 4, wherein the step of determining if an algorithm
2 required to process the at least one request for data mining processing is supported

3 comprises comparing the data defining at least one algorithm that must be performed
4 in order to perform the requested data mining processing with data defining
5 algorithms that are supported by the computer system.

1 6. The method of claim 5, wherein the data defining at least one algorithm that
2 must be performed in order to perform the requested data mining processing and the
3 data defining algorithms that are supported by the computer system are in extensible
4 markup language format.

1 7. The method of claim 5, wherein the step of determining whether the
2 computer system is available for additional processing comprises the step of:
3 determining whether available idle time of the computer system is greater
4 than a predefined or a dynamically calculated threshold.

1 8. The method of claim 5, wherein the computer system comprises a plurality of
2 processors and the step of determining whether the computer system is available for
3 additional processing comprises the step of:
4 determining whether any of the plurality of processors is available for additional
5 processing.

1 9. The method of claim 8, wherein the step of determining whether any of the
2 plurality of processors is available for additional processing comprises the step of
3 determining, for each of the plurality of processors, whether available idle
4 time of the processor is greater than a predefined or a dynamically calculated
5 threshold.

1 10. The method of claim 5, wherein the step of determining whether the
2 computer system is available for additional processing comprises the step of:
3 determining availability of the computer system for additional processing
4 relative to at least one other computer system.

1 11. The method of claim 5, wherein the step of determining whether the
2 computer system will become available for additional processing before other
3 computer systems that might process the at least one request comprises the steps of:
4 estimating a time to availability of the computer system;
5 exchanging an estimate of a time to availability of the at least one other
6 computer system; and
7 comparing the time to availability of the computer system with the time to
8 availability of the at least one other computer system.

1 12. The method of claim 11, wherein the step of determining whether the
2 computer system will be able to complete requested processing in an allotted time
3 comprises the steps of:

4 estimating a time to completion for the requested processing on the computer
5 system;

6 comparing the time to completion for the requested processing on the
7 computer system with time allocation information included in the request for data
8 mining processing.

1 13. In a data mining agent executing in a computer system, a method of data
2 mining comprising the steps of:

3 determining that the computer system is overloaded;

4 querying at least one other computer system to determine whether the at least
5 one other computer system can complete a data mining processing task being
6 performed on the computer system faster than the computer system;

7 determining whether the at least one other computer system can complete the
8 data mining processing task being performed on the computer system faster than the
9 computer system; and

10 if the at least one other computer system can complete the data mining
11 processing task faster than the computer system, migrating the processing from the
12 computer system to the at least one other computer system.

1 14. The method of claim 13, wherein the migrating step comprises the steps of:
2 reserving the at least one other computer system for migration;
3 interrupting and checkpointing the data mining processing task on the
4 computer system; and
5 enqueueing a request to the at least one other computer system for continued
6 processing of the data mining processing task.

1 15. The method of claim 14, wherein the step of determining that the computer
2 system is overloaded comprises the step of:
3 determining that the computer system is overloaded if a utilization of a
4 processor of the computer system is greater than a predefined threshold for a
5 predefined time.

1 16. The method of claim 15, wherein the querying step comprises the step of:
2 generating an estimate of a time to complete the data mining processing task.

1 17. The method of claim 16, wherein the generating step comprises the steps of:
2 estimating an amount of processing that must be performed to complete the
3 data mining processing task;

4 estimating a processor utilization that will be available to process the data
5 mining processing task; and
6 estimating a time to complete the data mining processing task based on the
7 estimate of the amount of processing that must be performed, the estimate of
8 available processor utilization, and a speed of the processor.

1 18. The method of claim 17, wherein the querying step further comprises the step
2 of:

3 requesting information from the at least one other computer system, the
4 information including a speed of the at least one other computer system and an
5 estimate of processor utilization of the at least one other computer system.

1 19. The method of claim 18, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system comprises the
4 step of:

5 estimating a time to complete the data mining processing task for the at least
6 one other computer system based on the estimate of the amount of processing that
7 must be performed to complete the data mining processing task, the speed of the at
8 least one other computer system and the estimate of processor utilization of the at
9 least one other computer system.

1 20. The method of claim 19, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system further
4 comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;

8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 21. The method of claim 17, wherein the querying step further comprises the step
2 of:

3 transmitting to the at least one other computer system the estimate of the
4 amount of processing that must be performed to complete the data mining processing
5 task; and

6 receiving from the at least one other computer system an estimate of a time to
7 complete the data mining processing task for the at least one other computer system

1 22. The method of claim 21, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system further
4 comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;

8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 23. In a data mining agent executing in a computer system, a method of data
2 mining comprising the steps of:

3 determining that a processing load in the computer system is high relative to
4 at least one other computer system;

5 determining a remaining cost of completing processing of a data mining
6 processing task being processed by the computer system;
7 determining whether the at least one other computer system can complete
8 processing of the data mining processing task at a lower cost than the computer
9 system; and
10 if the at least one other computer system can complete processing of the data
11 mining processing task at a lower cost than the computer system, migrating
12 processing of the data mining processing task to the at least one computer system.

1 24. The method of claim 23, wherein the step of determining that a processing
2 load in the computer system is high relative to at least one other computer system
3 comprises the steps of:
4 determining a processor utilization of the computer system;
5 determining a processor utilization of the at least one other computer system;
6 and
7 determining that the processor utilization of the computer system is greater
8 than a predefined amount higher than the processor utilization of the at least one
9 other computer system.

1 25. The method of claim 24, wherein the remaining cost of completing
2 processing of a data mining processing task is determined based on a time to
3 complete processing of the data mining processing task.

1 26. The method of claim 24, wherein the remaining cost of completing
2 processing of a data mining processing task is determined based on a time to
3 complete processing of the data mining processing task and on additional factors,
4 including actual costs of use of the computer system.

1 27. The method of claim 24, wherein the step of determining a remaining cost of
2 completing processing of a data mining processing task being processed by the
3 computer system comprises the steps of:

4 estimating an amount of processing that must be performed to complete the
5 data mining processing task;

6 estimating a processor utilization that will be available to process the data
7 mining processing task; and

8 estimating a time to complete the data mining processing task based on the
9 estimate of the amount of processing that must be performed, the estimate of
10 available processor utilization, and a speed of the processor.

1 28. The method of claim 27, further comprising the step of:

2 estimating additional factors, including actual costs of use of the computer
3 system.

1 29. The method of claim 27, wherein the step of determining whether the at least
2 one other computer system can complete processing of the data mining processing
3 task at a lower cost than the computer system comprises the step of:
4 soliciting a bid for completing processing of the data mining processing task
5 from the at least one other computer system.

1 30. The method of claim 29, wherein the soliciting step comprises the steps of:
2 transmitting a request for a bid to the at least one other computer system, the
3 request for the bid including information relating to the amount of processing that
4 must be performed to complete the data mining processing task; and
5 receiving a bid from the at least one other computer system, the bid including
6 an estimate of a cost of completing processing of the data mining processing task on
7 the at least one other computer system.

1 31. A computer system for performing data mining, comprising:
2 a processor operable to execute computer program instructions;
3 a memory operable to store computer program instructions executable by
4 the processor; and

5 a data mining agent comprising computer program instructions, for
6 performing the steps of:

7 examining a request queue comprising at least one request for data mining
8 processing,

9 determining if the at least one request for data mining processing can be
10 processed,

11 accepting the at least one request for data mining processing if it is
12 determined that the at least one request for data mining processing can be processed,
13 and

14 processing the accepted request for data mining processing in the computer
15 system.

1 32. The system of claim 31, wherein the determining step comprises the steps of:

2 determining if an algorithm required to process the at least one request for
3 data mining processing is supported by the computer system;

4 if the algorithm required to process the at least one request for data mining
5 processing is supported, determining whether the computer system is available for
6 additional processing;

7 if the computer system is not available for additional processing, determining
8 whether the computer system will become available for additional processing before
9 other computer systems that might process the at least one request;

10 if the computer system is available for additional processing, or if the
11 computer system will become available for additional processing before other
12 computer systems that might process the at least one request, determining whether
13 the computer system will be able to complete requested processing in an allotted
14 time; and

15 if the computer system will be able to complete the requested processing in
16 the allotted time, determining that the computer system can process the at least one
17 request for data mining processing.

1 33. The system of claim 32, wherein the at least one request for data mining
2 processing comprises data defining at least one algorithm that must be performed in
3 order to perform the requested data mining processing.

1 34. The system of claim 33, wherein there is data defining algorithms that are
2 supported by the computer system.

1 35. The system of claim 34, wherein the step of determining if an algorithm
2 required to process the at least one request for data mining processing is supported
3 comprises comparing the data defining at least one algorithm that must be performed
4 in order to perform the requested data mining processing with data defining
5 algorithms that are supported by the computer system.

1 36. The system of claim 35, wherein the data defining at least one algorithm that
2 must be performed in order to perform the requested data mining processing and the
3 data defining algorithms that are supported by the computer system are in extensible
4 markup language format.

1 37. The system of claim 35, wherein the step of determining whether the
2 computer system is available for additional processing comprises the step of:
3 determining whether available idle time of the computer system is greater
4 than a predefined or a dynamically calculated threshold.

1 38. The system of claim 35, wherein the computer system comprises a plurality of
2 processors and the step of determining whether the computer system is available for
3 additional processing comprises the step of:
4 determining whether any of the plurality of processors is available for additional
5 processing.

1 39. The system of claim 38, wherein the step of determining whether any of the
2 plurality of processors is available for additional processing comprises the step of
3 determining, for each of the plurality of processors, whether available idle

4 time of the processor is greater than a predefined or a dynamically calculated
5 threshold.

1 40. The system of claim 35, wherein the step of determining whether the
2 computer system is available for additional processing comprises the step of:

3 determining availability of the computer system for additional processing
4 relative to at least one other computer system.

1 41. The system of claim 35, wherein the step of determining whether the
2 computer system will become available for additional processing before other
3 computer systems that might process the at least one request comprises the steps of:

4 estimating a time to availability of the computer system;

5 exchanging an estimate of a time to availability of the at least one other
6 computer system; and

7 comparing the time to availability of the computer system with the time to
8 availability of the at least one other computer system.

1 42. The system of claim 41, wherein the step of determining whether the
2 computer system will be able to complete requested processing in an allotted time
3 comprises the steps of:

4 estimating a time to completion for the requested processing on the computer
5 system;
6 comparing the time to completion for the requested processing on the
7 computer system with time allocation information included in the request for data
8 mining processing.

1 43. A computer system for performing data mining, comprising:

2 a processor operable to execute computer program instructions;

3 a memory operable to store computer program instructions executable by
4 the processor; and

5 a data mining agent comprising computer program instructions, for
6 performing the steps of:

7 determining that the computer system is overloaded,

8 querying at least one other computer system to determine whether the at least
9 one other computer system can complete a data mining processing task being
10 performed on the computer system faster than the computer system,

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system, and

14 if the at least one other computer system can complete the data mining
15 processing task faster than the computer system, migrating the processing from the
16 computer system to the at least one other computer system.

1 44. The system of claim 43, wherein the migrating step comprises the steps of:
2 reserving the at least one other computer system for migration;
3 interrupting and checkpointing the data mining processing task on the
4 computer system; and
5 enqueueing a request to the at least one other computer system for continued
6 processing of the data mining processing task.

1 45. The system of claim 44, wherein the step of determining that the computer
2 system is overloaded comprises the step of:
3 determining that the computer system is overloaded if a utilization of a
4 processor of the computer system is greater than a predefined threshold for a
5 predefined time.

1 46. The system of claim 45, wherein the querying step comprises the step of:
2 generating an estimate of a time to complete the data mining processing task.

1 47. The system of claim 46, wherein the generating step comprises the steps of:

2 estimating an amount of processing that must be performed to complete the
3 data mining processing task;
4 estimating a processor utilization that will be available to process the data
5 mining processing task; and
6 estimating a time to complete the data mining processing task based on the
7 estimate of the amount of processing that must be performed, the estimate of
8 available processor utilization, and a speed of the processor.

1 48. The system of claim 47, wherein the querying step further comprises the step
2 of:

3 requesting information from the at least one other computer system, the
4 information including a speed of the at least one other computer system and an
5 estimate of processor utilization of the at least one other computer system.

1 49. The system of claim 48, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system comprises the
4 step of:

5 estimating a time to complete the data mining processing task for the at least
6 one other computer system based on the estimate of the amount of processing that
7 must be performed to complete the data mining processing task, the speed of the at

8 least one other computer system and the estimate of processor utilization of the at
9 least one other computer system.

1 50. The system of claim 49, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system further
4 comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;

8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 51. The system of claim 47, wherein the querying step further comprises the step
2 of:

3 transmitting to the at least one other computer system the estimate of the
4 amount of processing that must be performed to complete the data mining processing
5 task; and

6 receiving from the at least one other computer system an estimate of a time to
7 complete the data mining processing task for the at least one other computer system

1 52. The system of claim 51, wherein the step of determining whether the at least
2 one other computer system can complete a data mining processing task being
3 performed on the computer system faster than the computer system further
4 comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;

8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 53. A computer system for performing data mining, comprising:

2 a processor operable to execute computer program instructions;
3 a memory operable to store computer program instructions executable by
4 the processor; and
5 a data mining agent comprising computer program instructions, for
6 performing the steps of:
7 determining that a processing load in the computer system is high relative to
8 at least one other computer system,
9 determining a remaining cost of completing processing of a data mining
10 processing task being processed by the computer system,
11 determining whether the at least one other computer system can complete
12 processing of the data mining processing task at a lower cost than the computer
13 system, and
14 if the at least one other computer system can complete processing of the data
15 mining processing task at a lower cost than the computer system, migrating
16 processing of the data mining processing task to the at least one computer system.

1 54. The system of claim 53, wherein the step of determining that a processing
2 load in the computer system is high relative to at least one other computer system
3 comprises the steps of:
4 determining a processor utilization of the computer system;

5 determining a processor utilization of the at least one other computer system;
6 and
7 determining that the processor utilization of the computer system is greater
8 than a predefined amount higher than the processor utilization of the at least one
9 other computer system.

1 55. The system of claim 54, wherein the remaining cost of completing processing
2 of a data mining processing task is determined based on a time to complete
3 processing of the data mining processing task.

1 56. The system of claim 54, wherein the remaining cost of completing processing
2 of a data mining processing task is determined based on a time to complete
3 processing of the data mining processing task and on additional factors, including
4 actual costs of use of the computer system.

1 57. The system of claim 54, wherein the step of determining a remaining cost of
2 completing processing of a data mining processing task being processed by the
3 computer system comprises the steps of:
4 estimating an amount of processing that must be performed to complete the
5 data mining processing task;

6 estimating a processor utilization that will be available to process the data
7 mining processing task; and
8 estimating a time to complete the data mining processing task based on the
9 estimate of the amount of processing that must be performed, the estimate of
10 available processor utilization, and a speed of the processor.

1 58. The system of claim 57, further comprising the step of:
2 estimating additional factors, including actual costs of use of the computer
3 system.

1 59. The system of claim 57, wherein the step of determining whether the at least
2 one other computer system can complete processing of the data mining processing
3 task at a lower cost than the computer system comprises the step of:
4 soliciting a bid for completing processing of the data mining processing task
5 from the at least one other computer system.

1 60. The system of claim 59, wherein the soliciting step comprises the steps of:
2 transmitting a request for a bid to the at least one other computer system, the
3 request for the bid including information relating to the amount of processing that
4 must be performed to complete the data mining processing task; and

5 receiving a bid from the at least one other computer system, the bid including
6 an estimate of a cost of completing processing of the data mining processing task on
7 the at least one other computer system.

1 61. A computer program product for performing data mining in a computer
2 system data processing system, comprising:

3 a computer readable medium;

4 computer program instructions, recorded on the computer readable medium,
5 executable by a processor, for performing the steps of:

6 examining a request queue comprising at least one request for data mining
7 processing;

8 determining if the at least one request for data mining processing can be
9 processed;

10 accepting the at least one request for data mining processing if it is
11 determined that the at least one request for data mining processing can be processed;

12 and

13 processing the accepted request for data mining processing in the computer
14 system.

1 62. The computer program product of claim 61, wherein the determining step
2 comprises the steps of:

3 determining if an algorithm required to process the at least one request for
4 data mining processing is supported by the computer system;

5 if the algorithm required to process the at least one request for data mining
6 processing is supported, determining whether the computer system is available for
7 additional processing;

8 if the computer system is not available for additional processing, determining
9 whether the computer system will become available for additional processing before
10 other computer systems that might process the at least one request;

11 if the computer system is available for additional processing, or if the
12 computer system will become available for additional processing before other
13 computer systems that might process the at least one request, determining whether
14 the computer system will be able to complete requested processing in an allotted
15 time; and

16 if the computer system will be able to complete the requested processing in
17 the allotted time, determining that the computer system can process the at least one
18 request for data mining processing.

1 63. The computer program product of claim 62, wherein the at least one request
2 for data mining processing comprises data defining at least one algorithm that must
3 be performed in order to perform the requested data mining processing.

1 64. The computer program product of claim 63, wherein there is data defining
2 algorithms that are supported by the computer system.

1 65. The computer program product of claim 64, wherein the step of determining
2 if an algorithm required to process the at least one request for data mining processing
3 is supported comprises comparing the data defining at least one algorithm that must
4 be performed in order to perform the requested data mining processing with data
5 defining algorithms that are supported by the computer system.

1 66. The computer program product of claim 65, wherein the data defining at least
2 one algorithm that must be performed in order to perform the requested data mining
3 processing and the data defining algorithms that are supported by the computer
4 system are in extensible markup language format.

1 67. The computer program product of claim 65, wherein the step of determining
2 whether the computer system is available for additional processing comprises the
3 step of:
4 determining whether available idle time of the computer system is greater
5 than a predefined or a dynamically calculated threshold.

1 68. The computer program product of claim 65, wherein the computer system
2 comprises a plurality of processors and the step of determining whether the computer
3 system is available for additional processing comprises the step of:
4 determining whether any of the plurality of processors is available for additional
5 processing.

1 69. The computer program product of claim 68, wherein the step of determining
2 whether any of the plurality of processors is available for additional processing
3 comprises the step of
4 determining, for each of the plurality of processors, whether available idle
5 time of the processor is greater than a predefined or a dynamically calculated
6 threshold.

1 70. The computer program product of claim 65, wherein the step of determining
2 whether the computer system is available for additional processing comprises the
3 step of:
4 determining availability of the computer system for additional processing
5 relative to at least one other computer system.

1 71. The computer program product of claim 65, wherein the step of determining
2 whether the computer system will become available for additional processing before

3 other computer systems that might process the at least one request comprises the
4 steps of:

5 estimating a time to availability of the computer system;

6 exchanging an estimate of a time to availability of the at least one other
7 computer system; and

8 comparing the time to availability of the computer system with the time to
9 availability of the at least one other computer system.

1 72. The computer program product of claim 71, wherein the step of determining
2 whether the computer system will be able to complete requested processing in an
3 allotted time comprises the steps of:

4 estimating a time to completion for the requested processing on the computer
5 system;

6 comparing the time to completion for the requested processing on the
7 computer system with time allocation information included in the request for data
8 mining processing.

1 73. A computer program product for performing data mining in a computer
2 system data processing system, comprising:

3 a computer readable medium;

4 computer program instructions, recorded on the computer readable medium,
5 executable by a processor, for performing the steps of:
6 determining that the computer system is overloaded;
7 querying at least one other computer system to determine whether the at least
8 one other computer system can complete a data mining processing task being
9 performed on the computer system faster than the computer system;
10 determining whether the at least one other computer system can complete the
11 data mining processing task being performed on the computer system faster than the
12 computer system; and
13 if the at least one other computer system can complete the data mining
14 processing task faster than the computer system, migrating the processing from the
15 computer system to the at least one other computer system.

1 74. The computer program product of claim 73, wherein the migrating step
2 comprises the steps of:
3 reserving the at least one other computer system for migration;
4 interrupting and checkpointing the data mining processing task on the
5 computer system; and
6 enqueueing a request to the at least one other computer system for continued
7 processing of the data mining processing task.

1 75. The computer program product of claim 74, wherein the step of determining
2 that the computer system is overloaded comprises the step of:
3 determining that the computer system is overloaded if a utilization of a
4 processor of the computer system is greater than a predefined threshold for a
5 predefined time.

1 76. The computer program product of claim 75, wherein the querying step
2 comprises the step of:
3 generating an estimate of a time to complete the data mining processing task.

1 77. The computer program product of claim 76, wherein the generating step
2 comprises the steps of:

3 estimating an amount of processing that must be performed to complete the
4 data mining processing task;

5 estimating a processor utilization that will be available to process the data
6 mining processing task; and

7 estimating a time to complete the data mining processing task based on the
8 estimate of the amount of processing that must be performed, the estimate of
9 available processor utilization, and a speed of the processor.

1 78. The computer program product of claim 77, wherein the querying step further
2 comprises the step of:
3 requesting information from the at least one other computer system, the
4 information including a speed of the at least one other computer system and an
5 estimate of processor utilization of the at least one other computer system.

1 79. The computer program product of claim 78, wherein the step of determining
2 whether the at least one other computer system can complete a data mining
3 processing task being performed on the computer system faster than the computer
4 system comprises the step of:
5 estimating a time to complete the data mining processing task for the at least
6 one other computer system based on the estimate of the amount of processing that
7 must be performed to complete the data mining processing task, the speed of the at
8 least one other computer system and the estimate of processor utilization of the at
9 least one other computer system.

1 80. The computer program product of claim 79, wherein the step of determining
2 whether the at least one other computer system can complete a data mining
3 processing task being performed on the computer system faster than the computer
4 system further comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;
8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and
11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 81. The computer program product of claim 77, wherein the querying step further
2 comprises the step of:

3 transmitting to the at least one other computer system the estimate of the
4 amount of processing that must be performed to complete the data mining processing
5 task; and

6 receiving from the at least one other computer system an estimate of a time to
7 complete the data mining processing task for the at least one other computer system

1 82. The computer program product of claim 81, wherein the step of determining
2 whether the at least one other computer system can complete a data mining

3 processing task being performed on the computer system faster than the computer
4 system further comprises the steps of:

5 adding an estimate of a time to migrate the data mining processing task to the
6 at least one other computer system and the estimate of the time to complete the data
7 mining processing task for the at least one other computer system;

8 comparing the estimate of the time to complete the data mining processing
9 task for the computer system with the estimate of the time to complete the data
10 mining processing task for the at least one other computer system; and

11 determining whether the at least one other computer system can complete the
12 data mining processing task being performed on the computer system faster than the
13 computer system.

1 83. A computer program product for performing data mining in a computer
2 system data processing system, comprising:

3 a computer readable medium;

4 computer program instructions, recorded on the computer readable medium,
5 executable by a processor, for performing the steps of:

6 determining that a processing load in the computer system is high relative to
7 at least one other computer system;

8 determining a remaining cost of completing processing of a data mining
9 processing task being processed by the computer system;

10 determining whether the at least one other computer system can complete
11 processing of the data mining processing task at a lower cost than the computer
12 system; and

13 if the at least one other computer system can complete processing of the data
14 mining processing task at a lower cost than the computer system, migrating
15 processing of the data mining processing task to the at least one computer system.

1 84. The computer program product of claim 83, wherein the step of determining
2 that a processing load in the computer system is high relative to at least one other
3 computer system comprises the steps of:

4 determining a processor utilization of the computer system;

5 determining a processor utilization of the at least one other computer system;

6 and

7 determining that the processor utilization of the computer system is greater
8 than a predefined amount higher than the processor utilization of the at least one
9 other computer system.

1 85. The computer program product of claim 84, wherein the remaining cost of
2 completing processing of a data mining processing task is determined based on a
3 time to complete processing of the data mining processing task.

1 86. The computer program product of claim 84, wherein the remaining cost of
2 completing processing of a data mining processing task is determined based on a
3 time to complete processing of the data mining processing task and on additional
4 factors, including actual costs of use of the computer system.

1 87. The computer program product of claim 84, wherein the step of determining a
2 remaining cost of completing processing of a data mining processing task being
3 processed by the computer system comprises the steps of:

4 estimating an amount of processing that must be performed to complete the
5 data mining processing task;

6 estimating a processor utilization that will be available to process the data
7 mining processing task; and

8 estimating a time to complete the data mining processing task based on the
9 estimate of the amount of processing that must be performed, the estimate of
10 available processor utilization, and a speed of the processor.

1 88. The computer program product of claim 87, further comprising the step of:

2 estimating additional factors, including actual costs of use of the computer
3 system.

1 89. The computer program product of claim 87, wherein the step of determining
2 whether the at least one other computer system can complete processing of the data
3 mining processing task at a lower cost than the computer system comprises the step
4 of:

5 soliciting a bid for completing processing of the data mining processing task
6 from the at least one other computer system.

7 90. The computer program product of claim 89, wherein the soliciting step
8 comprises the steps of:

9 transmitting a request for a bid to the at least one other computer system, the
10 request for the bid including information relating to the amount of processing that
11 must be performed to complete the data mining processing task; and

12 receiving a bid from the at least one other computer system, the bid including
13 an estimate of a cost of completing processing of the data mining processing task on
14 the at least one other computer system.

1 91. A data mining agent for managing data mining in a computer system, the
2 data mining agent comprising:

3 means for examining a request queue comprising at least one request for data
4 mining processing;

5 means for determining if the at least one request for data mining processing
6 can be processed;

7 means for accepting the at least one request for data mining processing if it is
8 determined that the at least one request for data mining processing can be processed;
9 and

10 means for processing the accepted request for data mining processing in the
11 computer system.

1 92. The data mining agent of claim 91, wherein the determining means
2 comprises:

3 means for determining if an algorithm required to process the at least one
4 request for data mining processing is supported by the computer system;

5 means for determining whether the computer system is available for
6 additional processing, if the algorithm required to process the at least one request for
7 data mining processing is supported;

8 means for determining whether the computer system will become available
9 for additional processing before other computer systems that might process the at
10 least one request, if the computer system is not available for additional processing;

11 means for determining whether the computer system will be able to complete
12 requested processing in an allotted time, if the computer system is available for
13 additional processing, or if the computer system will become available for additional

14 processing before other computer systems that might process the at least one request;
15 and
16 means for determining that the computer system can process the at least one
17 request for data mining processing, if the computer system will be able to complete
18 the requested processing in the allotted time.

1 93. The data mining agent of claim 92, wherein the at least one request for data
2 mining processing comprises data defining at least one algorithm that must be
3 performed in order to perform the requested data mining processing.

1 94. The data mining agent of claim 93, wherein there is data defining algorithms
2 that are supported by the computer system.

1 95. The data mining agent of claim 94, wherein the means for determining if an
2 algorithm required to process the at least one request for data mining processing is
3 supported comprises means for comparing the data defining at least one algorithm
4 that must be performed in order to perform the requested data mining processing
5 with data defining algorithms that are supported by the computer system.

1 96. The data mining agent of claim 95, wherein the data defining at least one
2 algorithm that must be performed in order to perform the requested data mining

3 processing and the data defining algorithms that are supported by the computer
4 system are in extensible markup language format.

1 97. The data mining agent of claim 95, wherein the means for determining
2 whether the computer system is available for additional processing comprises:

3 means for determining whether available idle time of the computer system is
4 greater than a predefined or a dynamically calculated threshold.

1 98. The data mining agent of claim 95, wherein the computer system comprises a
2 plurality of processors and the means for determining whether the computer system
3 is available for additional processing comprises:

4 means for determining whether any of the plurality of processors is available
5 for additional processing.

1 99. The data mining agent of claim 98, wherein the means for determining
2 whether any of the plurality of processors is available for additional processing
3 comprises:

4 means for determining, for each of the plurality of processors, whether
5 available idle time of the processor is greater than a predefined or a dynamically
6 calculated threshold.

1 100. The data mining agent of claim 95, wherein the means for determining
 2 whether the computer system is available for additional processing comprises:
 3 means for determining availability of the computer system for additional
 4 processing relative to at least one other computer system.

1 101. The data mining agent of claim 95, wherein the means for determining
 2 whether the computer system will become available for additional processing before
 3 other computer systems that might process the at least one request comprises:
 4 means for estimating a time to availability of the computer system;
 5 means for exchanging an estimate of a time to availability of the at least one
 6 other computer system; and
 7 means for comparing the time to availability of the computer system with the
 8 time to availability of the at least one other computer system.

1 102. The data mining agent of claim 101, wherein the means for determining
 2 whether the computer system will be able to complete requested processing in an
 3 allotted time comprises:
 4 means for estimating a time to completion for the requested processing on the
 5 computer system;

6 means for comparing the time to completion for the requested processing on
7 the computer system with time allocation information included in the request for data
8 mining processing.

1 103. A data mining agent for managing data mining in a computer system, the
2 data mining agent comprising:

3 means for determining that the computer system is overloaded;

4 means for querying at least one other computer system to determine whether
5 the at least one other computer system can complete a data mining processing task
6 being performed on the computer system faster than the computer system;

7 means for determining whether the at least one other computer system can
8 complete the data mining processing task being performed on the computer system
9 faster than the computer system; and

10 means for migrating the processing from the computer system to the at least
11 one other computer system, if the at least one other computer system can complete
12 the data mining processing task faster than the computer system.

1 104. The data mining agent of claim 103, wherein the migrating means comprises:

2 means for reserving the at least one other computer system for migration;

3 means for interrupting and checkpointing the data mining processing task on
4 the computer system; and

5 means for enqueueing a request to the at least one other computer system for
6 continued processing of the data mining processing task.

1 105. The data mining agent of claim 104, wherein the means for determining that
2 the computer system is overloaded comprises:

3 means for determining that the computer system is overloaded if a utilization
4 of a processor of the computer system is greater than a predefined threshold for a
5 predefined time.

1 106. The data mining agent of claim 105, wherein the querying means comprises:

2 means for generating an estimate of a time to complete the data mining
3 processing task.

1 107. The data mining agent of claim 106, wherein the generating means
2 comprises:

3 means for estimating an amount of processing that must be performed to
4 complete the data mining processing task;

5 means for estimating a processor utilization that will be available to process
6 the data mining processing task; and

7 means for estimating a time to complete the data mining processing task
8 based on the estimate of the amount of processing that must be performed, the
9 estimate of available processor utilization, and a speed of the processor.

1 108. The data mining agent of claim 107, wherein the querying means further
2 comprises:

3 means for requesting information from the at least one other computer
4 system, the information including a speed of the at least one other computer system
5 and an estimate of processor utilization of the at least one other computer system.

1 109. The data mining agent of claim 18, wherein the means for determining
2 whether the at least one other computer system can complete a data mining
3 processing task being performed on the computer system faster than the computer
4 system comprises:

5 means for estimating a time to complete the data mining processing task for
6 the at least one other computer system based on the estimate of the amount of
7 processing that must be performed to complete the data mining processing task, the
8 speed of the at least one other computer system and the estimate of processor
9 utilization of the at least one other computer system.

1 110. The data mining agent of claim 19, wherein the means for determining
2 whether the at least one other computer system can complete a data mining
3 processing task being performed on the computer system faster than the computer
4 system further comprises:

5 means for adding an estimate of a time to migrate the data mining processing
6 task to the at least one other computer system and the estimate of the time to
7 complete the data mining processing task for the at least one other computer system;

8 means for comparing the estimate of the time to complete the data mining
9 processing task for the computer system with the estimate of the time to complete the
10 data mining processing task for the at least one other computer system; and

11 means for determining whether the at least one other computer system can
12 complete the data mining processing task being performed on the computer system
13 faster than the computer system.

1 111. The data mining agent of claim 17, wherein the querying means further
2 comprises:

3 means for transmitting to the at least one other computer system the estimate
4 of the amount of processing that must be performed to complete the data mining
5 processing task; and

6 means for receiving from the at least one other computer system an estimate
7 of a time to complete the data mining processing task for the at least one other
8 computer system

1 112. The data mining agent of claim 21, wherein the means for determining
2 whether the at least one other computer system can complete a data mining
3 processing task being performed on the computer system faster than the computer
4 system further comprises:

5 means for adding an estimate of a time to migrate the data mining processing
6 task to the at least one other computer system and the estimate of the time to
7 complete the data mining processing task for the at least one other computer system;

8 means for comparing the estimate of the time to complete the data mining
9 processing task for the computer system with the estimate of the time to complete the
10 data mining processing task for the at least one other computer system; and

11 means for determining whether the at least one other computer system can
12 complete the data mining processing task being performed on the computer system
13 faster than the computer system.

1 113. A data mining agent for managing data mining in a computer system, the
2 data mining agent comprising:

3 means for determining that a processing load in the computer system is high
4 relative to at least one other computer system;

5 means for determining a remaining cost of completing processing of a data
6 mining processing task being processed by the computer system;

7 means for determining whether the at least one other computer system can
8 complete processing of the data mining processing task at a lower cost than the
9 computer system; and

10 means for migrating processing of the data mining processing task to the at
11 least one computer system, the at least one other computer system can complete
12 processing of the data mining processing task at a lower cost than the computer
13 system.

1 114. The data mining agent of claim 113 wherein the step of determining that a
2 processing load in the computer system is high relative to at least one other computer
3 system comprises:

4 means for determining a processor utilization of the computer system;

5 means for determining a processor utilization of the at least one other
6 computer system; and

7 means for determining that the processor utilization of the computer system is
8 greater than a predefined amount higher than the processor utilization of the at least
9 one other computer system.

1 115. The data mining agent of claim 114, wherein the remaining cost of
2 completing processing of a data mining processing task is determined based on a
3 time to complete processing of the data mining processing task.

1 116. The data mining agent of claim 114, wherein the remaining cost of
2 completing processing of a data mining processing task is determined based on a
3 time to complete processing of the data mining processing task and on additional
4 factors, including actual costs of use of the computer system.

1 117. The data mining agent of claim 114, wherein the means for determining a
2 remaining cost of completing processing of a data mining processing task being
3 processed by the computer system comprises:

4 means for estimating an amount of processing that must be performed to
5 complete the data mining processing task;

6 means for estimating a processor utilization that will be available to process
7 the data mining processing task; and

8 means for estimating a time to complete the data mining processing task
9 based on the estimate of the amount of processing that must be performed, the
10 estimate of available processor utilization, and a speed of the processor.

1 118. The data mining agent of claim 117, further comprising:
2 means for estimating additional factors, including actual costs of use of the
3 computer system.

1 119. The data mining agent of claim 117, wherein the means for determining
2 whether the at least one other computer system can complete processing of the data
3 mining processing task at a lower cost than the computer system comprises:
4 means for soliciting a bid for completing processing of the data mining
5 processing task from the at least one other computer system.

1 120. The data mining agent of claim 119, wherein the soliciting means comprises:
2 means for transmitting a request for a bid to the at least one other computer
3 system, the request for the bid including information relating to the amount of
4 processing that must be performed to complete the data mining processing task; and
5 means for receiving a bid from the at least one other computer system, the bid
6 including an estimate of a cost of completing processing of the data mining
7 processing task on the at least one other computer system.